

MATROID FILTRATIONS AND COMPUTATIONAL PERSISTENT HOMOLOGY

GREGORY HENSELMAN AND ROBERT GHRIST

ABSTRACT. This technical report introduces a novel approach to efficient computation in homological algebra over fields, with particular emphasis on computing the persistent homology of a filtered topological cell complex. The algorithms here presented rely on a novel relationship between discrete Morse theory, matroid theory, and classical matrix factorizations. We provide background, detail the algorithms, and benchmark the software implementation in the EIRENE package.

1. INTRODUCTION

This paper reports recent work toward the goal of establishing a set of efficient computational tools in homological algebra over a field [13]. Our principal motivation is to improve upon the foundation built in [6] and develop and implement efficient algorithms for computing cosheaf homology and sheaf cohomology in the cellular setting. One can specialize to constructible cosheaf homology over the real line, which includes persistent homology as a particular instance. As this incidental application is of particular current relevance to topological data analysis [5, 8, 10, 11, 19, 22], we frame our methods to this setting in this technical report.

We present a novel algorithm for computing persistent homology of a filtered complex quickly and memory-efficiently. There are three ingredients in our approach.

- (1) **Matroid theory.** Our algorithms are couched in the language of matroids, both for elegance and efficiency. This theory permits simple reformulation of standard notions of persistence [filtrations, barcodes] in a manner that incorporates matroid concepts [rank, modularity, minimal bases] and the concomitant matroid algorithms.
- (2) **Discrete Morse theory.** It is well-known [2, 3, 7, 15, 16] that discrete Morse theory can be used to perform a simplification of the complex while preserving homology. We build on this to use the Morse theory to avoid construction and storage of the full boundary operator, thus improving memory performance.

- (3) **Matrix factorization.** Combining matroid and Morse perspectives with classical matrix factorization theorems allows for both more efficient computation as well as the ability to back-out representative generators of homology as chains in the *original* input complex, a capability of significant interest in topological data analysis.

This technical report is intended for a reader familiar with homology, persistent homology, computational homology, and (elementary) discrete Morse theory: as such, we pass over the usual literature review and basic background. Since (in the computational topology community) matroid theory is much less familiar, we provide enough background in §2 to explain the language used in our algorithms. Matrix reformulation of persistence and subsequent algorithms are presented in §3–§6. Implementation and benchmarking appears in §7.

2. BACKGROUND

For the purpose of continuity it will be convenient to fix a ground field \mathbb{k} , to be referenced throughout. Given a \mathbb{k} -vector space V and a set inclusion $S \hookrightarrow V$, we write $\langle S \rangle$ for the subspace spanned by S and V/S for the quotient of V by $\langle S \rangle$. Where context leaves no room for confusion we will identify the elements of V with their images under quotient maps, making it possible, for example, to speak unambiguously of the linear independence of T in V/S . Given an arbitrary set E and singleton $\{j\}$, we write $E + j$ for $E \cup \{j\}$, and given any binary relation R on E we write R^{op} for $\{(t, s) : (s, t) \in R\}$.

We review the elements of finite matroid theory, full details of which may be found in a number of excellent introductory texts [4, 20, 21, 25]. The reader should note that while, for historical reasons, most authors restrict scope of the term *matroid* to that of finite matroids, a number of generalizations to the infinite setting have been proposed. Among these the *finitary matroids of finite rank* suit our purposes best, and we refer to finite-rank, finitary matroids implicitly wherever matroids appear in the text.

With that preface, a (finite-rank, finitary) *matroid* M consists of a pair (E, \mathcal{I}) , where E is a set, $\mathcal{I} \subseteq 2^E$ is a family of subsets, and (i) $J \in \mathcal{I}$ whenever $I \in \mathcal{I}$ and $J \subseteq I$, and (ii) if $I, J \in \mathcal{I}$ with $|I| < |J|$, then there exists $j \in J - I$ such that $I + j \in \mathcal{I}$. The sets E and \mathcal{I} , sometimes written $E(M)$ and $\mathcal{I}(M)$ to emphasize association with M , are called the *ground set* and *independence system*, respectively. The elements of \mathcal{I} are *independent sets* and those of $2^E - \mathcal{I}$ are *dependent*. Maximal-under-inclusion independent sets are *bases* and minimal-under-inclusion dependent sets are *circuits*. It follows directly from the axioms that a singleton $\{e\}$ forms a circuit if and only if the replacement set

$$R(e \rightarrow B) = \{b \in B : B - b + e \in \mathcal{B}(M)\}$$

is empty for some basis B . The family of all bases is denoted $\mathcal{B}(M)$, and that of all circuits is denoted $\mathcal{C}(M)$.

The conditions imposed by (ii) on $\mathcal{I}(M)$, $\mathcal{B}(M)$, and $\mathcal{C}(M)$ support a remarkable branch of combinatorics, with fundamental ties to discrete and algebraic geometry, optimization, algebraic and combinatorial topology, graph theory, and analysis

of algorithms. They also underlie a number of the useful properties associated with finite-dimensional vector spaces, many of which, including dimension, have matroid-theoretic analogs.

To illustrate, consider a matroid found very often in literature: that of a subset U in a finite-dimensional, \mathbb{k} -linear vector space V . To every such U corresponds an independence system \mathcal{I} consisting of the \mathbb{k} -linearly independent subsets of U , and it is an elementary exercise to show that the pair $M(U) = (U, \mathcal{I})$ satisfies the axioms of a finite-rank, finitary matroid. We will introduce the next group of terms with explicit reference to this construction.

Zero We say $e \in E$ is a *zero element* if $\{e\} \in \mathcal{C}(M)$. As remarked above, the zero elements are those for which the set $R(e \rightarrow B)$ is empty. Clearly $u \in U$ is a zero in $M(U)$ if and only if $u = 0 \in V$.

Rank Axiom (ii) can be shown to imply that $|I| = |J|$ for every pair of maximal-with-respect-to-inclusion independent subsets of an arbitrary $S \subseteq E$. This number, denoted $\rho(S)$, is called the *rank* of S . By convention $\rho(M) = \rho(E)$. If $M = M(U)$, then $\rho(S) = \dim \langle S \rangle$.

Closure The *closure* of $S \subseteq E$, written $cl(S)$, is the discrete analog of a subspace generated by S in V : formally, \bar{S} is set of all $e \in E$ such that $\{e\} = C - S$ for some $C \in \mathcal{C}(M)$. If $M = M(U)$, then $\bar{S} = U \cap \langle S \rangle$. A *flat* is a subset that contains its closure.

Deletion If $S \subseteq E$ then the family $\mathcal{I} = \{T \in \mathcal{I}(M) : T \subseteq E - S\}$ is the independence system of a matroid $M - S$ on ground set $E - S$, called the *deletion* of M by S . The *restriction* of M to S , written $M|S$, is the deletion of M by $E - S$. If $M = M(U)$ then $M|S = M(S)$.

Contraction The *contraction* of M by S , denoted M/S , is the matroid on ground set $E - S$ with independence system $\{T \in \mathcal{I}(M) : T \cup J \in \mathcal{I} \ \forall \ J \in \mathcal{I} \cap 2^S\}$. For the purpose of our discussion it is worth noting an alternate characterization of $\mathcal{I}(M/S)$: if I is a maximal independent subset of S , then $\mathcal{I}(M/S) = \{J \subseteq E - S : J \cup I \in \mathcal{I}(M)\}$. As an immediate corollary, $\rho(M/S) = \rho(M) - \rho(S)$. If $M = M(U)$ then M/S is the matroid generated by $U - S$ in the quotient space $V/\langle S \rangle$.

Minor Matroids obtained by sequential deletion and contraction operations are *minors* of M . It can be shown that deletion and contraction commute, so that every minor may be expressed $(M - S)/T$ for some S and T . Where context leaves no room for confusion we write S/T for $(M|S)/T$.

Representation A (\mathbb{k} -linear) representation of M is a map $\phi : E \rightarrow \mathbb{k}^r$ such that $I \in \mathcal{I}$ if and only if $\phi(I)$ is linearly independent. Matroids that admit \mathbb{k} -linear representations are called \mathbb{k} -linear. It is common practice to identify a representation $E \rightarrow \mathbb{k}^r$ with the matrix $A \in \mathbb{k}^{\{1, \dots, r\} \times E}$ such that $A[:, e] = \phi(e)$ for all $e \in E$. In general,

we will not distinguish between A and the associated E -indexed family of column vectors. Given $B \in \mathcal{B}(M)$, we say φ has *B-standard form* if $\varphi(B)$ is the basis of standard unit vectors in \mathbb{k}^r . Clearly B -standard representations of $M(\mathbb{k}^r)$ are in 1-1 correspondence with $GL_r(\mathbb{k})$.

We note a few facts regarding discrete optimization. Let $f : E \rightarrow \mathbb{R}$ be any weight function, $\mathcal{S} \subseteq e^E$ be any family of subsets, and $\mathcal{B} \subseteq \mathcal{S}$ be the family of maximal-under-inclusion elements of \mathcal{S} . The elements of $\operatorname{argmax}_{B \in \mathcal{B}} f(B)$, where $f(B) = \sum_{b \in B} f(b)$, are of fundamental importance to a number of fields in discrete geometry and combinatorics, and the problem of finding the maximal B given E, f , and an oracle to determine membership in \mathcal{S} has been vigorously studied since the mid-20th century. This problem, NP-hard in general, admits a highly efficient greedy solution in the special case where (E, \mathcal{S}) constitutes a matroid. In fact, this characterizes finite matroids completely. The following is classical:

Proposition 1. *Let E be a finite set and \mathcal{S} be a subset of 2^E closed under inclusion. Then \mathcal{S} is the independence system of a matroid on ground set E if and only if Algorithm 1 returns an element of $\operatorname{argmax}_{\mathcal{S}} f$ for arbitrary $f : E \rightarrow \mathbb{R}$.*

Proposition 1 affords a number of efficient proof techniques in the theory of finite matroids (as will be seen). As $\operatorname{argmin}_{\mathcal{S}} f = \operatorname{argmax}_{\mathcal{S}} (-f)$, Algorithm 1 also provides a greedy method for obtaining minimal bases, which will be of primary concern.

Algorithm 1 Greedy algorithm for maximal set-weight

```

1: Label the elements of  $E$  by  $e_1, \dots, e_n$ , such that  $f(e_{i+1}) \leq f(e_i)$ .
2:  $S := \emptyset$ 
3: for  $i = 1, \dots, n$  do
4:   if  $S + e_i \in \mathcal{S}$  then
5:      $S \leftarrow S + e_i$ 
6:   end if
7: end for
```

Our final comments concern the relationship between a pair of f -minimal bases. Up to this point, the terms and observations introduced have been standard elements of matroid canon. In the following sections we will need two new notions, namely those of a flat function and an f -triangular matrix.

Given $S, T \subseteq E$ and $f : E \rightarrow \mathbb{Z}$, say that $A \in \mathbb{k}^{S \times T}$ is *f-upper triangular* if $f[s] \leq f[t]$ whenever $A[s, t]$ is nonzero. The product of two f -upper triangular matrices is again upper-triangular, as are their inverses. We say A is *f-lower triangular* if it is g -upper triangular for $g = -f$, and *f-diagonal* if it is f -upper and f -lower triangular. We say f is *flat* if the inverse image $f^{-1}(-\infty, t]$ is a flat of M for every choice of t .

Lemma 2. *Let $f : E \rightarrow \mathbb{Z}$ be a flat weight function; B, F bases of M ; and R_B the unique element of $\{0, 1\}^{B \times E}$ such that*

$$R_B[b, e] = \begin{cases} 1 & \text{if } B - b + e \in \mathcal{B}(M) \\ 0 & \text{otherwise.} \end{cases}$$

If B is f -minimal, then F is f -minimal if and only if $R_B[B, F]$ is f -upper triangular.

Lemma 3. *Let M be a linear matroid and f a flat weight function on M . If F is an f -minimal basis and $L \in GL_F(\mathbb{k})$, $R \in GL_E(\mathbb{k})$ are f -upper triangular then a basis B is f -minimal if and only if $LA[F, B]R$ is f -upper triangular for every F -standard representation A .*

Lemma 4. *Let f be a flat weight function on $M(\mathbb{k}^E)$, and for each $S \subseteq E$ let A_S be the unique element of $\mathbb{k}^{E \times S}$ such that $A[E, s] = s$ for $s \in S$. If S is finite and $R \in GL_S(\mathbb{k})$, $L \in GL_E(\mathbb{k})$ are f -upper triangular, then A_S contains an f -minimal basis if and only if $LA_S R$ contains one, also.*

3. MODULAR FILTRATIONS

To make precise the place of linear persistence in discrete optimization, we introduce one further notion, that of modular filtration. Much work in the recent field of homological data analysis may be viewed as a treatment of the relationship between two or more filtrations on a vector space. By analogy, we define a filtration \mathcal{F} of matroid M to be a nested sequence of flats $\emptyset = \mathcal{F}_0 \subseteq \dots \subseteq \mathcal{F}_L = E$. The *characteristic function* of \mathcal{F} is the map $\chi_{\mathcal{F}} : E \rightarrow \mathbb{Z}_{\geq 0}$ such that $\mathcal{F}_k = \chi_{\mathcal{F}}^{-1}\{0, \dots, k\}$ for all k . A pair of filtrations $(\mathcal{F}, \mathcal{G})$ is *modular* if

$$(1) \quad \rho(\mathcal{F}_i \cap \mathcal{G}_j) + \rho(\mathcal{G}_i \cup \mathcal{F}_j) = \rho(\mathcal{F}_i) + \rho(\mathcal{G}_j)$$

for all i and j . Modularity is of marked structural significance in the general theory of matroids, and offers a powerful array of tools for the analysis of filtrations.

Proposition 5. *If \mathcal{F}, \mathcal{G} are filtrations of M , then the following are equivalent.*

- (1) $(\mathcal{F}, \mathcal{G})$ is modular.
- (2) There exists a basis of minimal weight with respect to both $\chi_{\mathcal{F}}$ and $\chi_{\mathcal{G}}$.
- (3) There exists a basis B such that

$$B \cap (\mathcal{F}_i \cap \mathcal{G}_j) \in \mathcal{B}(\mathcal{F}_i \cap \mathcal{G}_j)$$

for all i, j .

The proof is organized as follows. Equivalence of (2) and (3) follows from Lemma 6 below.¹ It is readily seen that (1) follows from (3), as the latter implies $\rho(\mathcal{F}_i \cap \mathcal{G}_j) = |B \cap \mathcal{F}_i \cap \mathcal{G}_j|$, thereby reducing the identity that defines modularity to

$$|B \cap \mathcal{F}_i \cap \mathcal{G}_j| = |B \cap \mathcal{F}_i| + |B \cap \mathcal{G}_j| - |B \cap \mathcal{F}_i \cap \mathcal{G}_j|.$$

It remains to be shown, therefore, that (1) implies either (2) or (3). Given our special interest in minimal bases, a constructive proof of the former is to be desired.

¹Lemma 6 yields, moreover, a highly efficient means of checking independence over large families of intersections, and will be used implicitly throughout.

Lemma 6. *A basis B has minimal weight with respect to $\chi_{\mathcal{F}}$ and $\chi_{\mathcal{G}}$ if and only if*

$$(2) \quad B \cap (\mathcal{F}_i \cap \mathcal{G}_j) \in \mathcal{B}(\mathcal{F}_i \cap \mathcal{G}_j)$$

for all i, j .

Proof. If B satisfies (2) for all i, j then minimality with respect to \mathcal{F} follows from an application of Lemma 6 to the family of intersections $B \cap \mathcal{F}_i \cap E$. Minimality with respect to \mathcal{B} follows likewise. If on the other hand $S = B \cap (\mathcal{F}_i \cap \mathcal{G}_j) \notin \mathcal{B}(\mathcal{F}_i \cap \mathcal{G}_j)$ for some i, j , then there exists $s \in \mathcal{F}_i \cap \mathcal{G}_j$ such that $S + s \in \mathcal{I}(M)$. The fundamental circuit of s with respect to B intersects $B - S$ nontrivially, hence $B - b + s \in \mathcal{B}(M)$ for some $b \in B - S$. Either $\chi_{\mathcal{F}}(s) < \chi_{\mathcal{F}}(b)$ or $\chi_{\mathcal{G}}(s) < \chi_{\mathcal{G}}(b)$, and the desired conclusion follows. \square

Proof of Proposition 5. As already discussed, it suffices to show (1) implies (2). Therefore assume $(\mathcal{F}, \mathcal{G})$ is modular, and for each i fix a $\chi_{\mathcal{G}}$ -minimal

$$B_i \in \mathcal{B}(\mathcal{F}_i / \mathcal{F}_{i-1}).$$

The union $B = \cup_i B_i \in \mathcal{B}(M)$ forms a $\chi_{\mathcal{F}}$ -minimal basis in M , so it suffices to show B is minimal with respect to \mathcal{G} . Since $|B \cap \mathcal{G}_j| \leq \rho(\mathcal{G}_j)$, we may do so by proving $|B \cap \mathcal{G}_j| \geq \rho(\mathcal{G}_j)$ for all j . Therefore let i and j be given, fix $S_{i-1} \in \mathcal{B}(\mathcal{F}_{i-1} \cap \mathcal{G}_j)$ and extend S_{i-1} to a basis S_i of $\mathcal{F}_i \cap \mathcal{G}_j$.

Modularity of $(\mathcal{F}_i, \mathcal{G}_j)$ implies

$$\rho(\mathcal{G}_j / \mathcal{F}_i) = \rho(\mathcal{G}_j \cup \mathcal{F}_i) - \rho(\mathcal{F}_i) = \rho(\mathcal{G}_j) - \rho(\mathcal{G}_j \cap \mathcal{F}_i) = \rho(\mathcal{G}_j / (\mathcal{G}_j \cap \mathcal{F}_i)),$$

so that, for all i and j ,

$$\rho(\mathcal{G}_j / \mathcal{F}_i) = \rho(\mathcal{G}_j / (\mathcal{F}_i \cap \mathcal{G}_j)) = \rho(\mathcal{G}_j / S_i) \leq \rho(\mathcal{G}_j / (S_i \cup \mathcal{F}_{i-1})) \leq \rho(\mathcal{G}_j / \mathcal{F}_i).$$

As the quantities on either side are identical, strict equality holds throughout. Thus the second equality below.

$$\begin{aligned} \rho(\mathcal{G}_j / S_{i-1}) - \rho(S_i / S_{i-1}) &= \rho(\mathcal{G}_j / S_i) \\ &= \rho(\mathcal{G}_j / (S_i \cup \mathcal{F}_{i-1})) = \rho(\mathcal{G}_j / \mathcal{F}_{i-1}) - \rho(S_i / \mathcal{F}_{i-1}). \end{aligned}$$

A comparison of left- and right-hand sides shows $\rho(S_i / S_{i-1}) = \rho(S_i / \mathcal{F}_{i-1})$. By construction the set $T_{ij} = S_i - S_{i-1} \in \mathcal{G}_j \cap \mathcal{F}_i - \mathcal{F}_{i-1}$ forms a basis in S_i / S_{i-1} , hence

$$|T_{ij}| = \rho(S_i / S_{i-1}) = \rho(S_i / \mathcal{F}_{i-1}) = \rho(T_{ij} / \mathcal{F}_{i-1}).$$

Thus $T_{ij} \in \mathcal{I}(\mathcal{F}_i / \mathcal{F}_{i-1})$, so $|B_i \cap \mathcal{G}_j| \geq |T_{ij}|$. A second and third application of modularity provide the second and third equalities below,

$$\begin{aligned} |T_{ij}| &= \rho(\mathcal{G}_j / S_{i-1}) - \rho(\mathcal{G}_j / S_i) \\ &= \rho(\mathcal{G}_j / \mathcal{F}_{i-1}) - \rho(\mathcal{G}_j / \mathcal{F}_i) \\ &= \rho(\mathcal{G}_j \cap \mathcal{F}_i) - \rho(\mathcal{G}_j \cap \mathcal{F}_{i-1}) \end{aligned}$$

whence $|B_i \cap \mathcal{G}_j| \geq \rho(\mathcal{G}_j \cap \mathcal{F}_i) - \rho(\mathcal{G}_j \cap \mathcal{F}_{i-1})$. Summing over i yields

$$|B \cap \mathcal{G}_j| \geq \rho(\mathcal{G}_j)$$

which was to be shown. \square

4. LINEAR FILTRATIONS

We now specialize to the case of linear filtrations. Recall that a linear filtration of a finite-dimensional vector space V is a nested sequence of subspaces $\emptyset = \mathcal{F}_0 \subseteq \dots \subseteq \mathcal{F}_L = V$. Clearly, the linear filtrations of V are exactly the matroid-theoretic filtrations of $M(V)$. Since $\dim \mathcal{F}_i + \dim \mathcal{G}_j = \dim \mathcal{F}_i \cap \mathcal{G}_j + \dim \langle \mathcal{F}_i \cup \mathcal{G}_j \rangle$ when \mathcal{F} and \mathcal{G} are linear, Proposition 5 implies that to every such pair corresponds a basis $B \in \mathcal{B}(M)$ such that $\mathcal{F}_i \cap \mathcal{G}_j = \langle B \cap \mathcal{F}_i \cap \mathcal{G}_j \rangle$ for all i, j . We are interested in computing B when $V = \mathbb{k}^r$.

Assume that the data available for this pursuit are (1) an \mathcal{F} -minimal basis F , (2) a finite set G containing a $\chi_{\mathcal{G}}$ -minimal basis, and (3) oracles to evaluate $\chi_{\mathcal{F}}$ on F and $\chi_{\mathcal{G}}$ on G . These resources are commonly accessible for computations in persistent homology, as described in the following section. One can assume, further, that F is the set of standard unit vectors in $V = \mathbb{k}^F$. If such is not the case one can solve the analogous problem for an F -standard representation $\varphi \in GL(V)$, and port the solution back along φ^{-1} . Our strategy will be to construct a $(\chi_{\mathcal{F}}, \chi_{\mathcal{G}})$ -minimal basis from linear combinations of the vectors in G . A few pieces of notation will aid its description.

Given a matrix $A \in \mathbb{k}^{F \times G}$, write $\text{supp } s$ and $\text{supp } t$ for the supports of row and column vectors indexed by s and t , respectively. Say that a relation R on E respects $f : E \rightarrow \mathbb{R}$, or that R is an f -relation, if $f(s) \leq f(t)$ whenever $(s, t) \in R$. Given a $\chi_{\mathcal{F}}$ -linear order $<_{\mathcal{F}}$ and a $\chi_{\mathcal{G}}$ -linear order $<_{\mathcal{G}}$, define \mathcal{P} , S , and T by

$$\begin{aligned} \mathcal{P} &= \{(f, g) : f = \max_{<_{\mathcal{F}}} \text{supp } g, g = \min_{<_{\mathcal{G}}} \text{supp } f\} \\ S &= \{f : (f, g) \in \mathcal{P}(A, <_{\mathcal{F}}, <_{\mathcal{G}}) \text{ for some } g\} \\ T &= \{g : (f, g) \in \mathcal{P}(A, <_{\mathcal{F}}, <_{\mathcal{G}}) \text{ for some } f\} \end{aligned}$$

and L , R , X , Y , Z , and $*$ by

$$L = \begin{array}{c} S \\ S^c \end{array} \begin{array}{|c|c|} \hline X^{-1} & 0 \\ \hline -X^{-1}Z & I \\ \hline \end{array} \quad A = \begin{array}{c} S \\ S^c \end{array} \begin{array}{|c|c|} \hline X & Y \\ \hline Z & * \\ \hline \end{array} \quad R = \begin{array}{c} T \\ T^c \end{array} \begin{array}{|c|c|} \hline I & -X^{-1}Y \\ \hline 0 & I \\ \hline \end{array}$$

where $S^c = F - S$ and $T^c = G - T$. It can be helpful to regard \mathcal{P} as the Pareto-optimal frontier of $\text{supp } A = \{(s, t) : A[s, t] \neq 0\}$ with respect to $<_{\mathcal{G}}$ and the inverse-order $<_{\mathcal{F}}^{\text{op}}$. See Figure 1.

Lemma 7 now holds by construction.

Lemma 7. *If \mathcal{P} , S , and T are as above, then*

- (1) L is $\chi_{\mathcal{F}}$ -lower-triangular;
- (2) R is $\chi_{\mathcal{G}}$ -upper-triangular; and
- (3) The product LAR satisfies

$$LAR = \begin{array}{c} S \\ S^c \end{array} \begin{array}{|c|c|} \hline I & 0 \\ \hline 0 & * \\ \hline \end{array}$$

for some block-submatrix $*$ of rank equal to $\text{rank } A - |S|$.

Where convenient we will write $L(A, \langle \mathcal{F}, \langle \mathcal{G} \rangle)$, $R(A, \langle \mathcal{F}, \langle \mathcal{G} \rangle)$, and $\mathcal{P}(A, \langle \mathcal{F}, \langle \mathcal{G} \rangle)$ for L, R , and \mathcal{P} to emphasize association with their genera. In the description of Algorithm 2 we will write L_t and R_t for values of L and R associated to A_t .

Algorithm 2 Matrix reduction

```

1: for  $t = 1, \dots, r$  do
2:    $A_{t+1} \leftarrow L_t A_t R_t$ 
3:   if  $|A_{t+1}|_\infty = r$  then
4:     break
5:   end if
6: end for

```

The proof of Lemma 8 follows easily from Lemma 7:

Lemma 8. *Stopping condition $|A_{t+1}|_\infty = r$ holds for some $1 \leq t \leq r$. The matrix products $L = L_t \cdots L_1$ and $R = R_1 \cdots R_t$ are $\chi_{\mathcal{F}}$ -upper-triangular and $\chi_{\mathcal{G}}$ -lower-triangular, respectively.*

Corollary 9. *The columns of L^{-1} form an $(\mathcal{F}, \mathcal{G})$ -minimal basis in $M(V)$.*

Proof. If L and R are as in Lemma 8, then up to permutation and nonzero-scaling $LAR = [I \mid 0]$. By Lemma 4, $AR = [L^{-1} \mid 0]$ contains a \mathcal{G} -minimal basis, perforce the set of column vectors of L^{-1} . Since in addition L^{-1} is $\chi_{\mathcal{F}}$ -upper-triangular, the associated basis is also $\chi_{\mathcal{F}}$ -minimal. \square

It can be shown that $L(A, \langle \mathcal{F}, \langle \mathcal{G} \rangle) = L(AR, \langle \mathcal{F}, \langle \mathcal{G} \rangle)$, where $R = R(A, \langle \mathcal{F}, \langle \mathcal{G} \rangle)$. Consequently, the matrices R_t in Algorithm 2 need not play any role whatever in the calculation of L^{-1} : one must simply modify the stopping criterion to accommodate the fact that LA may have more nonzero coefficients than LAR . See Algorithm 3.

Algorithm 3 Light matrix reduction

```

1: Initialize  $A_1 = A$ 
2: while  $|\mathcal{P}(A_t)| < r$  do
3:    $A_{t+1} \leftarrow L_t A_t$ ,  $t \leftarrow t + 1$ 
4: end while

```

5. HOMOLOGICAL PERSISTENCE

An application of the n th homology functor to a nested sequence of chain complexes $C = (C_0 \subseteq \cdots \subseteq C_L)$ in $\mathbb{k}\text{-Vect}$ yields a diagram of vector spaces

$$\cdots \rightarrow 0 \rightarrow H_n(C_0) \rightarrow \cdots \rightarrow H_n(C_L) \rightarrow 0 \rightarrow \cdots$$

to which one may associate a graded $\mathbb{k}[t]$ -module $\mathcal{H}_n(C) = \oplus_i H_n(C_i)$, with t -action inherited from the induced map $H_n(C_i) \rightarrow H_n(C_{i+1})$. Modules of the form $\mathcal{H}_n(C)$, often called *persistence modules* [5, 26], are of marked interest in topological data analysis.




FIGURE 1. *Left:* The sparsity pattern of a matrix A with rows and columns in ascending $<_{\mathcal{F}}$ and $<_{\mathcal{G}}$ order. Nonzero coefficients are shaded in light or dark grey, the latter marking $\mathcal{P}(A, <_{\mathcal{F}}, <_{\mathcal{G}}) = \{(1,1), (4,2), (5,6), (6,7), (9,9)\}$. *Right:* The sparsity pattern of LAR , for generic A .

The sequence $C_0 \subseteq \cdots \subseteq C_L$ engenders a modular pair $(\mathcal{F}, \mathcal{G})$ on the n th cycle space $Z_n(C_L)$, where

$$\mathcal{F}_i = \ker \partial_n(C_i); \quad \mathcal{G}_i = \text{im } \partial_n(C_i); \quad i = 1, \dots, L;$$

and $\mathcal{G}_{L+1} = Z_n(C_L)$. Each n -cycle z maps naturally to a homology class $[z] \in H_n(C_i)$, where $i = \min\{j : z \in \mathcal{F}_j\}$. The proof of Proposition 10 follows from commutativity of Figure 5.

Proposition 10. *A basis $Z \in \mathcal{B}(Z_n)$ is $(\chi_{\mathcal{F}}, \chi_{\mathcal{G}})$ -minimal if and only if the non-nullhomologous cycles, $\{[z] : z \in Z, [z] \neq 0\}$, freely generate $\mathcal{H}_n(C)$.*

Computation of $(\chi_{\mathcal{F}}, \chi_{\mathcal{G}})$ -optimal bases requires slightly more than a rote application of Algorithm 2 in general, since representations of the cycle space Z_n seldom present *a priori*. Rather, the starting data is generally an indexed family E_1, \dots, E_L of $\chi_{\mathcal{F}}$ -minimal bases for the chain groups of C_L in all dimensions. A linear $\chi_{\mathcal{F}}$ -order on $E = E_1 \cup \cdots \cup E_n$, denoted $<$, and a matrix representation of each boundary operator with respect to E_n , denoted A_n , constitute the input to Algorithm 4. As with Algorithm 3, we write $\mathcal{P}(A_n)$ for $\mathcal{P}(A_n, <, <)$ and L_n, R_n for $L(A_n, <, <)$ and $R(A_n, <, <)$ respectively. Where n is not specified, each declaration should be understood for all n .

Proposition 11. *Algorithm 4 terminates for some $t = t_0$. The columns of*

$$Z = L_{n+1}^{-1} R_n [E_n, E_{n+1} - P_{n+1}]$$

form a $(\chi_{\mathcal{F}}, \chi_{\mathcal{G}})$ -minimal basis of Z_n , where $P_{n+1} = \{s : (s, t) \in P_{n+1}^0\}$.

Proof. Since $L_{n+1}^{-1} R_n$ is $\chi_{\mathcal{F}}$ -upper triangular, Z is $\chi_{\mathcal{F}}$ -minimal in its closure, Z_n . By Lemma 4, the columns of $R_n^{-1} L_{n+1} \partial_{n+1} L_{n+2}^{-1} R_{n+1}$ contain a $\chi_{\mathcal{G}}$ -minimal basis

Algorithm 4 Chain reduction

```

1: Initialize  $\mathcal{P}_n^0 = \emptyset$ ,  $\mathcal{P}_n^1 = \mathcal{P}(A_n)$ ,  $t = 1$ 
2: while  $\mathcal{P}_n^t \neq \mathcal{P}_n^{t-1}$  for some  $n$  do
3:    $t \leftarrow t + 1$ 
4:    $A_n \leftarrow L_n A_n L_{n+1}^{-1}$ 
5:    $\mathcal{P}_n^t = \mathcal{P}(A_n)$ 
6: end while
7:  $A_n \leftarrow R_{n-1}^{-1} A R_n$ 

```

The Barcode of a Modular Pair




FIGURE 2. Visual representation of a $(\chi_{\mathcal{F}}, \chi_{\mathcal{G}})$ -minimal basis $B = \{b^1, \dots, b^r\}$ in the special case where $\mathcal{G}_i \subseteq \mathcal{F}_i$, $i = 1, 2, \dots$. Each row corresponds to an element of B . The dot at location (i, b) is black if $b \in \mathcal{G}_i$, blue if $b \in \mathcal{F}_i - \mathcal{G}_i$, and grey otherwise. The blue dots in column i thus collectively represent a basis of $\mathcal{F}_i/\mathcal{G}_i$, the black dots represent zero elements, and the grey represent elements not contained in \mathcal{F}_i . The horizontal bars generated by the blue dots correspond exactly to the barcode of $(\mathcal{F}, \mathcal{G})$, when \mathcal{F} and \mathcal{G} are the induced filtrations on C_n for some filtered chain complex C .

of $cl(\partial_{n+1}) = \mathcal{G}_L^n$, perforce the subset indexed by P_{n+1} . In the E_n -standard representation of C_n this basis corresponds to $L_{n+1}^{-1} R_n[E_n, P_{n+1}]$. Any basis containing $L_{n+1}^{-1} R_n[E_n, P_{n+1}]$, in particular Z , is therefore \mathcal{G} -minimal. \square

6. ACYCLIC RELATIONS

Consider a nested sequence of cellular spaces $X_0 \subseteq \dots \subseteq X_L = X$. Most such filtered complexes that arise in scientific applications have very large n -skeleta, even

when n is small. Their boundary operators are highly sparse, and become accessible to computation only through sparse representation in memory. One drawback of sparse representation, however, is an oft disproportionate increase in the cost of computing matrix products. All things being equal, it is therefore preferable to execute fewer iterations of Algorithm 4 than many when A is sparsely represented. Since work stops when $|P'_n| = \text{rank} A_n$, one might naively hope to see improved performance when P_n^1 is close to $\text{rank} A_n$ – a hope realized in practice.

Since P'_n is entirely determined by $<$, anyone looking to maximize $|P_n^1|$ will do so by an informed choice of linear order. Enter acyclic relations, which provide a means to determine *a priori* the inclusion of certain sets, called acyclic matchings, in P_n^1 . Our strategy will be to find a favorable acyclic matching, and from this to engineer a compatible order $<$. Much effort has already been invested in the design and application of matchings in algebra, topology, combinatorics, and computation via the discrete Morse Theory of R. Forman and subsequent literature [9, 14, 16, 17, 24]. We will discuss the details of our own approach to matchings, and its connection to discrete Morse theory in [13]. At present we limit ourselves to a description of its use in Algorithm 4.

A binary relation R on a ground set E is *acyclic* if the transitive closure of R is antisymmetric. Evidently, R is acyclic if and only if the transitive closure of $R \cup \Delta(E \times E)$ is a partial order on E . The following observation is similarly clear, but bears record for ease of reference.

Lemma 12. *If E is finite and f is a real-valued function on E , then every acyclic f -relation extends to an f -linear order.*

Proof. Assume for convenience $\text{im } f = \{1, \dots, n\}$. Fix an acyclic f -relation R and for $i = 1, \dots, n$ let Λ_i be a linear order on $f^{-1}(i)$ respecting the transitive closure of $R \cup \Delta(E \times E)$. The closure of

$$\{(s, t) : f(s) < f(t)\} \cup \Lambda_1 \cup \dots \cup \Lambda_n$$

is a linear order of the desired form. \square

Suppose now that K is a finite-dimensional, \mathbb{k} -linear chain complex supported on $\{1, \dots, N\}$, that E_n freely generates K_n , and that $\deg(e) = n$ for each $e \in E_n$. If A_n is the matrix representation of ∂_n with respect to this basis, then a *matching* on (K, E) is a subrelation

$$V \subseteq \text{supp } A_1 \cup \dots \cup \text{supp } A_N$$

such that for each s_0 and t_0 in E , the pairs (s_0, t) and (s, t_0) belong to V for at most one s and one t . “Flipping” V yields a relation

$$R_V = \leq_{\deg} - V \cup V^{op}$$

where $\leq_{\deg} = \{(s, t) : \deg(s) < \deg(t)\}$. We say V is *acyclic* (*f*-acyclic) if R_V is acyclic (*f*-acyclic). Taken together, Proposition 13 and Lemma 12 imply that every *f*-acyclic matching includes into $\mathcal{P}(A_n, <)$ for some *f*-linear $<$. In particular, given large V , one can always find large P_n^1 .

Proposition 13. *If V is a matching and $<^{\text{op}}$ linearizes R_V , then V includes into $\cup_n \mathcal{P}(A_n, <)$.*

It is quite easy to check whether V is acyclic in practice. The task of deciding f -acyclisity, though somewhat more involved in general, reduces to a linear search when f is the characteristic function of a filtration.

Lemma 14. *An acyclic matching V is $\chi_{\mathcal{F}}$ -acyclic if and only if $\chi_{\mathcal{F}}(s) = \chi_{\mathcal{F}}(t)$ for all $(s, t) \in V$.*

We close with a brief but computationally useful observation concerning the order of operations whereby products of form $L_n A L_{n+1}$ are computed in Algorithm 4. Elementary calculations show that $A_n L_{n+1}$ the columns of $A_n L_{n+1}$ indexed by P'_{n+1} vanish at each step of the process. Thus the problem of computing $L_n A_n L_{n+1}$ reduces to that of computing $L_n A [E_{n-1}, E_n - P'_n]$. In particular, if only free generators for the first $N - 1$ homology groups are desired, then while there is no need to compute P_m for $m > N$, identifying a subset S of P_{N+1} allows one to reduce the calculation of $L_N A_N$ to that of $L_N A_N [E_N, E_{N+1} - S]$.

Top-dimensional boundary operators have special status in scientific computation, as they are often the the largest by far. If S may be determined formulaically prior to the construction and storage of A_N – for example, by determining a closed-form expression for an acyclic matching – then the cost of generating large portions of this matrix may be avoided altogether. As reported in the Experiments section, the effect of this reduction may be to drop the memory-cost of computation by several orders of magnitude. To our knowledge, this is the first principled use of acyclic matchings to avoid the construction not only of large portions of the cellular boundary operator, but of the underlying complex itself.

7. EXPERIMENTS

An instance of Algorithm 4 incorporating the optimization described in §6 has been implemented in the EIRENE library for homological algebra [12]. Experiments were conducted on a personal computer with Intel Core i7 processor at 2.3GHz, with 4 cores, 6MB of L3 Cache, and 16GB of RAM. Each core has 256 KB of L2 Cache. Results for a number of sample spaces, including those appearing in recent published benchmarks, are reported in Tables 1, 2, and 3. All homologies are computed using \mathbb{Z}_2 coefficients.

Our first round of experiments computes persistent homology of a Vietoris-Rips filtration on a random point cloud on n vertices in \mathbb{R}^{20} , for values of n up to 240. Persistent homology with representative generators is computed in dimensions 1, 2, and 3, with the total elapsed time and memory load (heap) recorded in Table 1.

Our second round of experiments parallels the benchmarks published in Fall 2015 [18]. Note that Tables 6.1 and 6.2 of this reference record time and space expenditures for certain large point clouds on various publicly available software packages, some of which are run on a cluster. We append one new example to this table: RG1E4, a randomly generated point cloud of 10,000 points in \mathbb{R}^{20} . This complex, the 2-skeleton of which has over 160 billion simplices, is, at the time

# Vertices	Size (B)	Time (s)	Heap(GB)	CR
40	0.00	2.13	0.02	0.14
80	0.00	4.36	0.02	0.07
120	0.19	15.7	4.58	0.04
160	0.82	44.1	21.65	0.03
200	2.54	124	33.46	0.03
240	6.36	407	53.18	0.02

TABLE 1. Persistence with generators in dimensions 1, 2, and 3 for the Vietoris-Rips complex of point clouds sampled from the uniform distribution on the unit cube in \mathbb{R}^{20} . *Size* refers to the size of the 4-skeleton of the complete simplex on n vertices. *CR*, or compression ratio, is the quotient of the size of the generated subcomplex by the 4-skeleton of the underlying space.

VR Complex	Size (M)	Time (s)	Heap (GB)
C. elegans	4.37	1.33	0.00
Klein	10.1	1.76	0.01
HIV	214	12.6	0.04
Dragon 1	166	15.6	0.08
Dragon 2	1.3k	141	2.32
RG1E4	1.66k	3.12k	35.7

TABLE 2. One-dimensional persistence for various Vietoris-Rips complexes. The data for C. elegans, Klein, HIV, Dragon 1, and Dragon 2 were drawn directly from the published benchmarks in [18]. RG1E4 is a random geometric complex on 10^4 vertices, sampled from the uniform distribution on the unit cube in \mathbb{R}^{20} .

of this writing, the largest complex whose homology we have computed. All the instances in Table 2 record computation of persistent H_1 .

The third set of experiments shows where the algorithm encounters difficulties. We compute higher dimensional homology of fixed (non-filtered) complexes arising from combinatorics. These complexes – the *matching* and *chessboard* complexes – are notoriously difficult to work with, as they have very few vertices, very many higher-dimensional simplices, and relatively large homology [23]. The performance of EIRENE in Table 3 is consistent with the expected difficulty: acyclic compression can only do so much for such complexes.

Finally, Figure 4 illustrates the degree of compression achieved through EIRENE in the context of a random point cloud in dimension 20. The horizontal axis records the number of vertices, n , used in the complex. The vertical axis records a ratio

Complex	Dim	Size (M)	Time (s)	Heap (GB)
Chessboard	7	1.44	8.38k	26.9
Matching	3	0.42	9.98k	37.0

TABLE 3. Homology of two unfiltered spaces, the chessboard complex $C_{8,8}$ and the matching complex $M_{3,13}$.

of size of various quantities relative to the rank of the boundary operator ∂_4 of the complete simplex on n vertices.




FIGURE 3. Persistent H_3 for a family of random geometric complexes. Samples of cardinality k were drawn from the uniform distribution on the unit cube in \mathbb{R}^{20} , $k = 20, \dots, 240$. The method of §6 was applied to the distance matrix d of each sample, resulting in a morse complex M . Recall that the n -cells of M are indexed by $M_n = E_n - P_n^1 - P_{n+1}^1$, where E_n is the family of n -faces of the simplex on k vertices, and P^1 is an acyclic matching. The EIRENE library applies a dynamic construction subroutine to build M_n from d directly. This subroutine generates the elements of $X_n = E_n - P_{n-1}^1$ sequentially and stores the elements of M_n in memory; it does not generate elements of $E_n - X_n$, nor does it store any combinatorial n -simplex in memory. In the figure above, black, red, and blue correspond to the ratios of $|E_4|$, $|X_4|$, and $|M_4|$, respectively, to $\dim \partial_4(E_4)$. All statistics represent averages taken across 10 samples.



FIGURE 4. A one-dimensional class representative. Grey points represent a sample of 5×10^3 points drawn from a torus embedded in \mathbb{R}^3 , with uniform random noise. Free generators for the associated persistence module, thresholded at three times the maximum noise level, were computed with the EIRENE library for homological algebra. A representative for the unique 1-dimensional class that survived to infinity was plotted with the open-source visualization library *Plotly* [1]. Vertices incident to the cycle representative appear in black.

ACKNOWLEDGEMENTS

The authors wish to express their sincere gratitude to C. Giusti and V. Nanda for their encyclopedic knowledge, useful comments, and continued encouragement. This work is supported by US DoD contracts FA9550-12-1-0416, FA9550-14-1-0012, and NO0014-16-1-2010.

REFERENCES

- [1] Plotly technologies inc. collaborative data science. <https://plot.ly>, 2015.
- [2] BAUER, U. *Persistence in discrete Morse theory*. PhD thesis, Göttingen, 2011.
- [3] BAUER, U., KERBER, M., AND REININGHAUS, J. Distributed computation of persistent homology. *Proceedings of Algorithm Engineering and Experiments (ALENEX)* (2014).
- [4] BJÖRNER, A., VERGNAS, M. L., STURMFELS, B., WHITE, N., AND ZIEGLER, G. M. *Oriented Matroids*. Cambridge University Press, 1993.
- [5] CARLSSON, G. Topology and data. *Bull. Amer. Math. Soc. (N.S.)* 46, 2 (2009), 255–308.
- [6] CURRY, J., GHRIST, R., AND NANDA, V. Discrete Morse theory for computing cellular sheaf cohomology. *ArXiv e-prints* (Dec. 2013).

- [7] DŁOTKO, P., KACZYŃSKI, T., MROZEK, M., AND WANNER, T. Coreduction homology algorithm for regular CW-Complexes. *Discrete & Computational Geometry* 46, 2 (2011), 361–388.
- [8] EDELSBRUNNER, H., AND HARER, J. *Computational Topology: an Introduction*. American Mathematical Society, Providence, RI, 2010.
- [9] FORMAN, R. Morse theory for cell complexes. *Adv. Math.* 134, 1 (1998), 90–145.
- [10] GHRIST, R. *Elementary Applied Topology*. Createspace, 2014.
- [11] GIUSTI, C., PASTALKOVA, E., CURTO, C., AND ITSKOV, V. Clique topology reveals intrinsic geometric structure in neural correlations. *Proceedings of the National Academy of the Sciences* (2015).
- [12] HENSELMAN, G. Eirene: a platform for computational homological algebra. <http://gregoryhenselman.org/eirene.html>, May 2016.
- [13] HENSELMAN, G. *Matroids, Filtrations, and Applications*. PhD thesis, University of Pennsylvania, 2016.
- [14] KOZLOV, D. Discrete Morse theory for free chain complexes. *Comptes Rendus Mathematique* 340 (2005), 867–872.
- [15] KOZLOV, D. *Combinatorial Algebraic Topology*, vol. 21 of *Algorithms and Computation in Mathematics*. Springer, 2008.
- [16] MISCHAIKOW, K., AND NANDA, V. Morse theory for filtrations and efficient computation of persistent homology. *Discrete Comput. Geom.* 50, 2 (2013), 330–353.
- [17] NANDA, V., TAMAKI, D., AND TANAKA, K. Discrete Morse theory and classifying spaces. in preparation, 2013.
- [18] OTTER, N., PORTER, M. A., TILLMANN, U., GRINDROD, P., AND HARRINGTON, H. A. A roadmap for the computation of persistent homology. <http://arxiv.org/abs/1506.08903>.
- [19] OUDOT. *Persistence Theory: From Quiver Representations to Data Analysis*. AMS, 2015.
- [20] OXLEY, J. *Matroid Theory*, 2nd ed. ed. Oxford University Press, 2011.
- [21] PAPDIMITRIOU, C., AND STEIGLITZ, K. *Combinatorial Optimization: Algorithms and Complexity*. Dover Press, 1998.
- [22] ROBINSON, M. *Topological Signal Processing*. Springer, Heidelberg, 2014.
- [23] SHARESHIAN, J., AND WACHS, M. Torsion in the matching complex and chessboard complex. *Advances in Math* 212 (2007), 525–570.
- [24] SKÖLDBERG, E. Morse theory from an algebraic viewpoint. *Transactions of the American Mathematical Society* 358, 1 (2006), 115–129.
- [25] TRUEMPER, K. *Matroid Decomposition*. Academic Press, San Diego, 1992.
- [26] ZOMORODIAN, A., AND CARLSSON, G. Computing persistent homology. *Discrete Comput. Geom.* 33, 2 (2005), 249–274.

DEPARTMENT ELECTRICAL/SYSTEMS ENGINEERING, UNIVERSITY OF PENNSYLVANIA, PHILADELPHIA PA, USA

E-mail address: grh@seas.upenn.edu

DEPARTMENTS OF MATHEMATICS AND ELECTRICAL/SYSTEMS ENGINEERING, UNIVERSITY OF PENNSYLVANIA, PHILADELPHIA PA, USA

E-mail address: ghrist@math.upenn.edu




FIGURE 5. A commutative diagram in Set . All maps preserve zero-elements. Recall that $H = B \cup \{0\} \subseteq \mathbb{k}^r$ for some (χ_F, χ_G) -optimal basis B . Arrows between the top 12 objects are inclusions. The lowest two vertical arrows are quotient maps in \mathbb{k} -Vect. The restriction of an arrow $a : X \rightarrow Y$ to codomain $X - \{0\}$ is an inclusion of matroid bases when a is oblique and $X \subseteq H$; it is an inclusion of vector bases when a is the composition of colinear oblique arrows. The diagram may be extended arbitrarily far to the right.